

# Compiler-Assisted Object Inlining with Value Fields

—  
**Rodrigo Bruno**, INESC-ID / Técnico - ULisboa\*

**Vojin Jovanovic**, Oracle Labs Switzerland

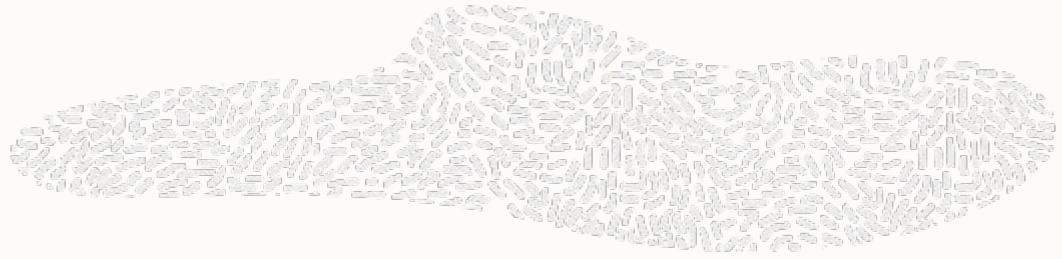
**Christian Wimmer**, Oracle Labs USA

**Gustavo Alonso**, Department of Computer Science, ETH Zurich

\* work done while at Oracle Labs Switzerland and ETH Zurich

# Data *objectification*

---

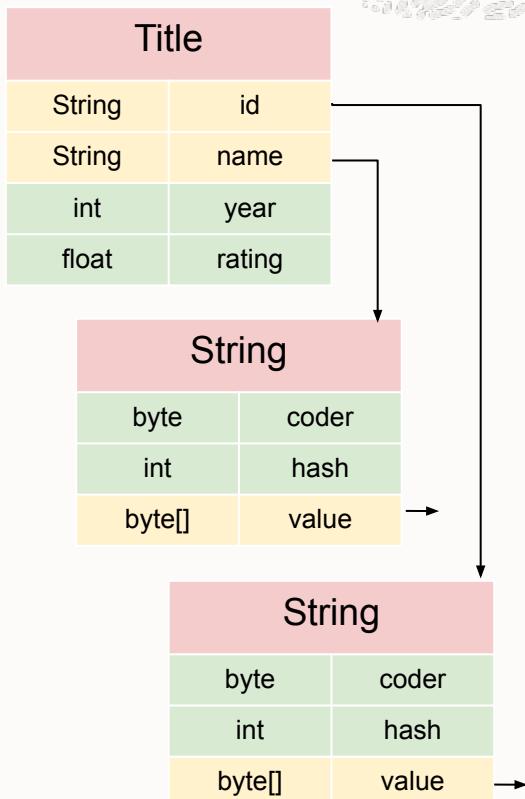


Title ID	Name	Year	Rating
...	...	...	...
tt0110912	Pulp Fiction	1994	8.9
...	...	...	...

Database  
(1 entry)

# Data objectification

Title ID	Name	Year	Rating
...	...	...	...
tt0110912	Pulp Fiction	1994	8.9
...	...	...	...



Database  
(1 entry)



Java Heap  
(5 objects, 4 references)

Data objectification

# Data objectification

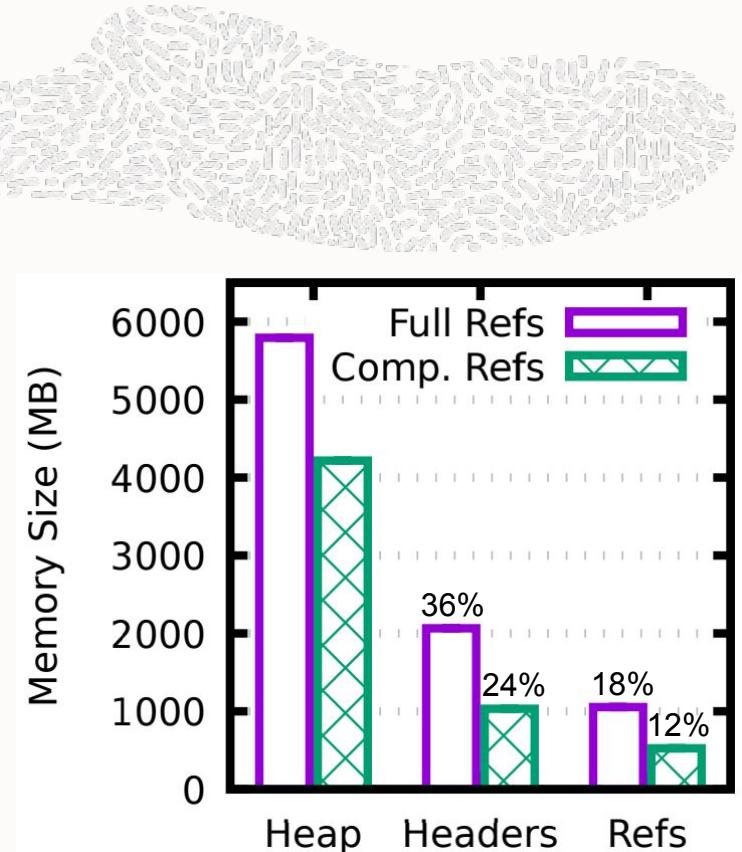
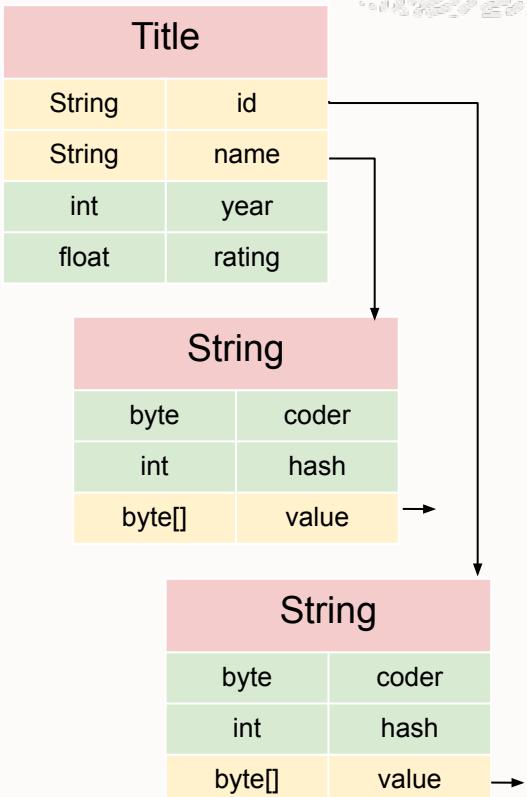
Title ID	Name	Year	Rating
...	...	...	...
tt0110912	Pulp Fiction	1994	8.9
...	...	...	...

Database  
(1 entry)

Data objectification



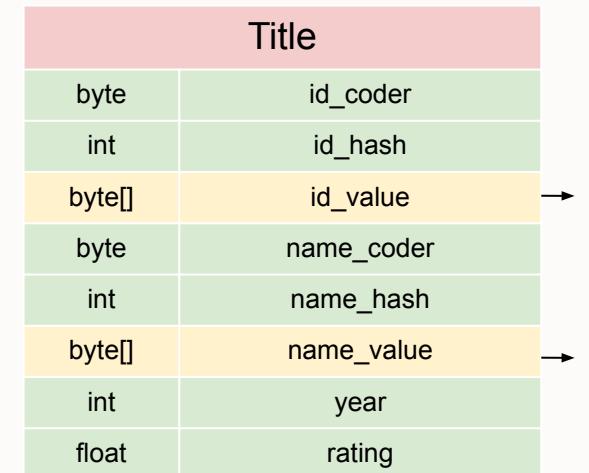
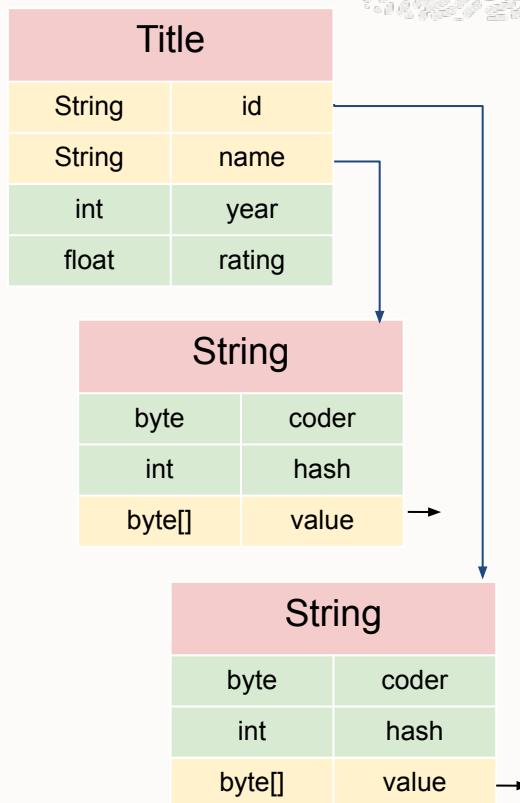
Java Heap  
(5 objects, 4 references)



'unusable' mem ~ 36-54%

# Data objectification

Title ID	Name	Year	Rating
...	...	...	...
tt0110912	Pulp Fiction	1994	8.9
...	...	...	...



Database  
(1 entry)

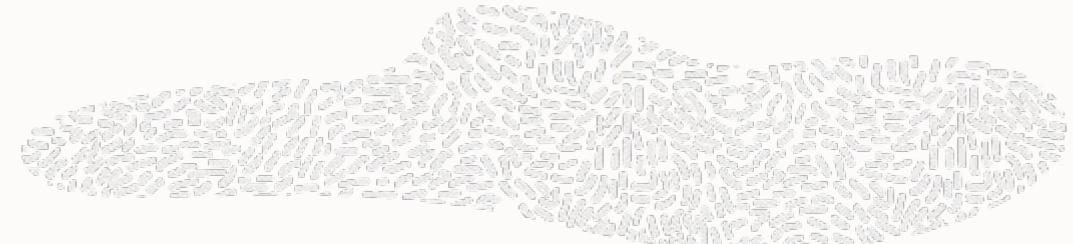
Java Heap  
(5 objects, 4 references)

Java Heap  
(3 objects, 2 references)

Data objectification

Object Inlining

# Value Fields



- **Simple abstraction:** fields with value semantics
  - opportunity for optimization (hint for the compiler)
    - fields have do not require identity
    - field accesses have do not require atomic reference access
- **Closed-world Assumption**
  - GraalVM Native Image
  - all types are known at compile time
- **Easy to use**
  - Field annotations: `@ValueField`, `@ValueGraph`
  - JSON configuration file
- **Target**
  - Java applications
  - allow framework/library developers to efficient data structures
  - end-users will automatically benefit

# Inlined Field Memory Layout

```
class Point {  
    final int x, y;  
}  
  
class Line {  
    @ValueField Point a, b;  
}
```

*child type*

*child fields*

*parent type*

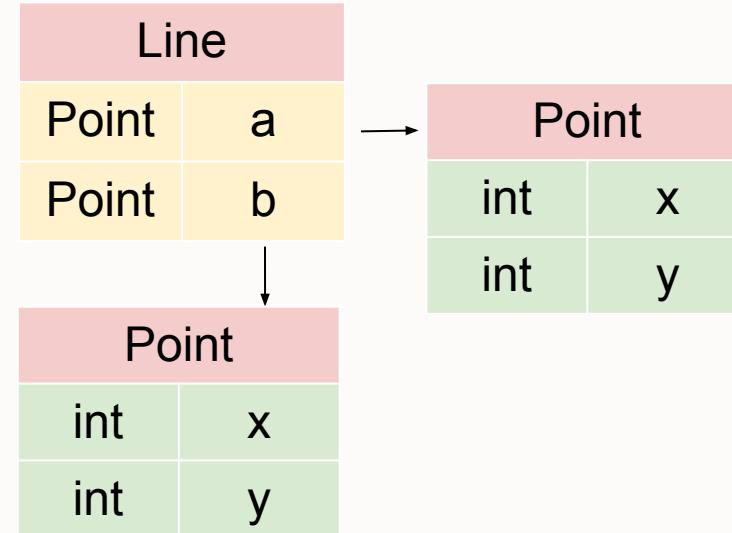
*parent fields*

Java source

# Inlined Field Memory Layout

```
class Point {  
    final int x, y;  
}  
  
class Line {  
    @ValueField Point a, b;  
}
```

*child type*  
*child fields*  
*parent type*  
*parent fields*



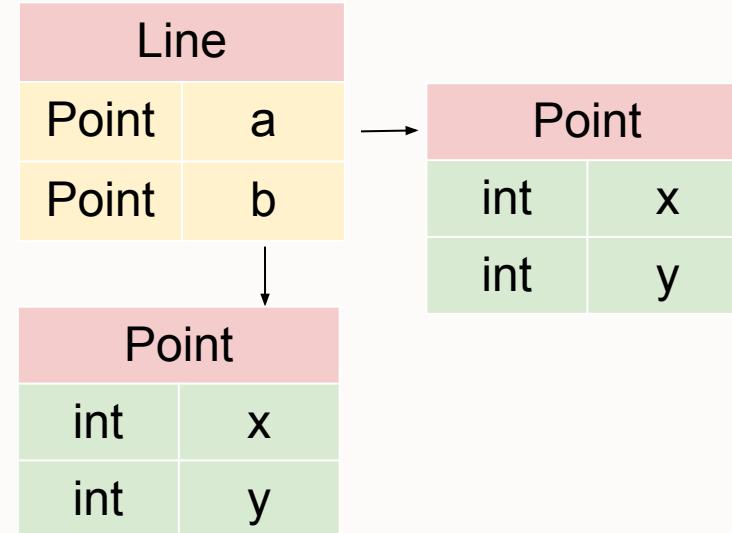
Java source

# Inlined Field Memory Layout

```
class Point {  
    final int x, y;  
}  
  
class Line {  
    @ValueField Point a, b;  
}
```

*child type*  
*child fields*  
*parent type*  
*parent fields*

Java source

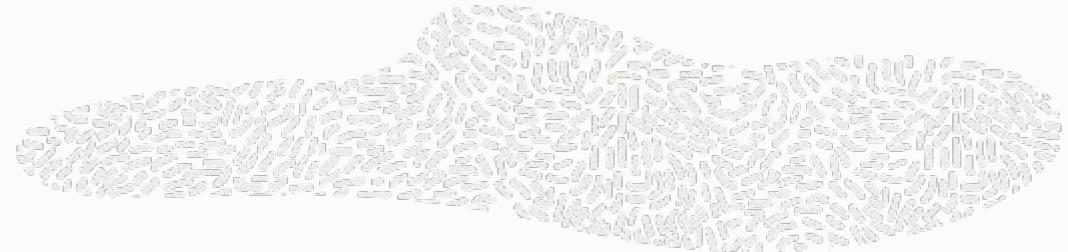


Original Object Layout

Line	
byte	a_state
int	a_x
int	a_y
byte	b_state
int	b_x
int	b_y

Inlined Object Layout

# Restrictions of Object Inlining

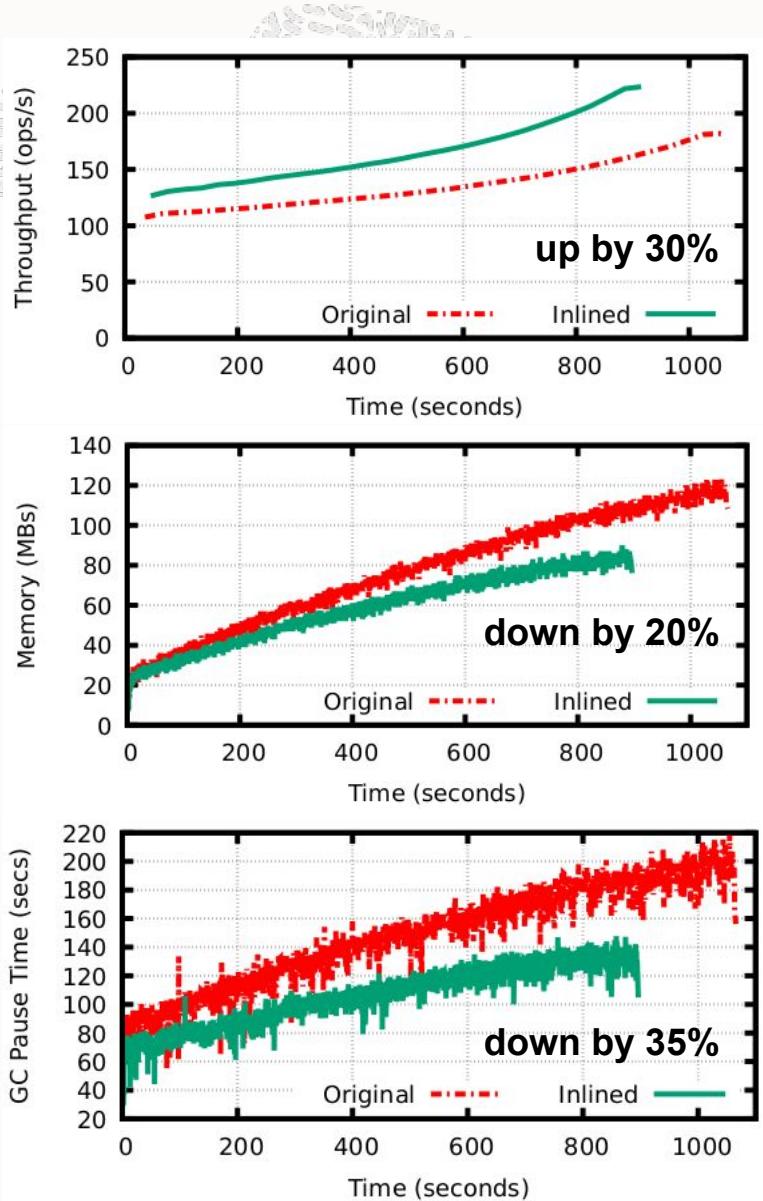


- Inlining is restricted to fields of **immutable** types
  - Modifications are not guaranteed to be propagated after inlining
  - If inlined objects are **not modified** after inlined, it is **safe** to inline **mutable** types
    - `@ValueField(allowMutable=true)`, `@ValueGraph(allowMutable=true)`
- No **polymorphic** types
  - We need to know the type layout at image build time
- No fields of **primitive** types, **array** type, **volatile**, **static**

# Spring Boot Online Website

- PetClinic workload:
  - Apache Jmeter produces requests
  - Microservice (Spring Boot / Micronaut) handles requests
  - SQL DB persists information
  - Realistic dataset (names of pets, real addresses, etc)
  - Workload consists of read and write requests
- Values are cached in memory (locally or remotely)
- Object inlining configuration file:

```
{  
    "value_fields" : {  
        "petclinic.model BaseEntity" : ["id"],  
        "petclinic.model Person" : ["firstName", "lastName"],  
        "petclinic.model NamedEntity" : ["name"],  
        "petclinic.owner Owner" : ["address", "city", "phone"],  
        "petclinic.owner Pet" : ["birthDate", "type"],  
        "petclinic.visit Visit" : ["date", "desc", "petId"]  
    }  
}
```

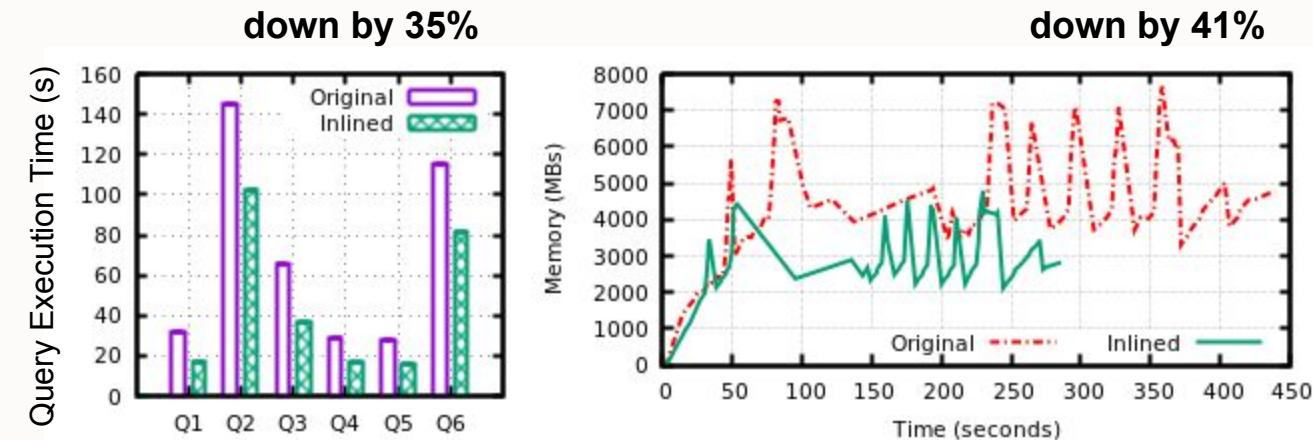


# Data Processing in Spark

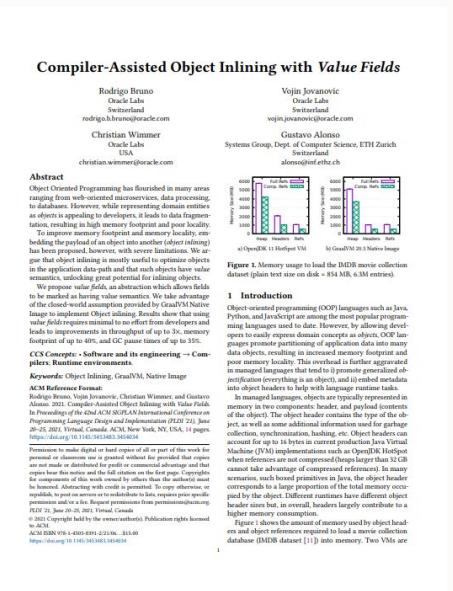
- Data analytics workload:
  - Large JSON dataset (IMDB movie collection)
  - 6 different queries are executed over the dataset
- Each JSON entry is loaded into a Spark RDD
  - Value Inlining can be used to inline fields inside simple POJO types

```
1: class Movie {           1: class Actor {  
2:   Actor[] actors;      2:   @ValueField Date birth;  
3:   @ValueField String genre; 3:   @ValueField Date death;  
4:   @ValueField String name; 4:   @ValueField String name;  
5:   @ValueField Date release; 5: }  
6:   int votes;  
7:   float rating;  
8: }
```

Query	Description
Q1	Number of movies released in a year by genre.
Q2	Movies ordered by movie rating.
Q3	Average age of a movie's actors.
Q4	Actors ordered by number of roles.
Q5	Year with more average rating vores.
Q6	Actors ordered by the number of roles in highly rated movies.



**More details available in the paper:**



# Rodrigo Bruno

[rodrigo.bruno@tecnico.ulisboa.pt](mailto:rodrigo.bruno@tecnico.ulisboa.pt)

<https://rodrigo-bruno.github.io/>

# Thank You



ORACLE